arXiv:1610.08148v1 [q-bio.PE] 26 Oct 2016

Online Bayesian phylogenetic inference: theoretical foundations via Sequential Monte Carlo

Vu Dinh[1] Aaron E. Darling[2] Frederick A. Matsen IV[1]

[1]Program in Computational Biology
Fred Hutchinson Cancer Research Center
Seattle, WA, USA

[2]The ithree institute
University of Technology Sydney
Ultimo NSW, Australia

Abstract

Phylogenetics, the inference of evolutionary trees from molecular sequence data such as DNA, is an enterprise that yields valuable evolutionary understanding of many biological systems. Bayesian phylogenetic algorithms, which approximate a posterior distribution on trees, have become a popular if computationally expensive means of doing phylogenetics. Modern data collection technologies are quickly adding new sequences to already substantial databases. With all current techniques for Bayesian phylogenetics, computation must start anew each time a sequence becomes available, making it costly to maintain an up-to-date estimate of a phylogenetic posterior. These considerations highlight the need for an *online* Bayesian phylogenetic method which can update an existing posterior with new sequences.

Here we provide theoretical results on the consistency and stability of methods for online Bayesian phylogenetic inference based on Sequential Monte Carlo (SMC) and Markov chain Monte Carlo (MCMC). We first show a consistency result, demonstrating that the method samples from the correct distribution in the limit of a large number of particles. Next we derive the first reported set of bounds on how phylogenetic likelihood surfaces change when new sequences are added. These bounds enable us to characterize the theoretical performance of sampling algorithms by bounding the effective sample size (ESS) with a given number of particles from below. We show that the ESS is guaranteed to grow linearly as the number of particles in an SMC sampler grows. Surprisingly, this result holds even though the dimensions of the phylogenetic model grow with each new added sequence.

MSC 2010 subject classifications: Primary 05C05, 60J22; secondary 92D15, 92B10.

Keywords: phylogenetics, Sequential Monte Carlo, effective sample size, online inference, Bayesian inference, subtree optimality

Funding: VD and FAM funded by National Science Foundation grants DMS-1223057 and CISE-1564137. FAM supported by a Faculty Scholar grant from the Howard Hughes Medical Institute and the Simons Foundation.

1 Background and main results

Phylogenetics is the theory and practice of reconstructing evolutionary trees. Evolutionary trees have found wide application in biology and medicine, including use in epidemiology, conservation planning, and cancer genomics. Maximum likelihood and Bayesian methods are generally considered to be the most powerful and accurate approaches for phylogenetic inference. The Bayesian methods in particular enjoy the flexibility to incorporate a wide range of ancillary model features such as geographical information or trait data which are essential for some applications. However, Bayesian tree inference with current implementations is a computationally intensive task, often requiring days or weeks of CPU time to analyze modest datasets with 100 or so sequences.

New developments in DNA and RNA sequencing technology have led to sustained growth in sequence datasets. This advanced technology has enabled real time outbreak surveillance efforts, such as ongoing Zika, Ebola, and foodborne disease sequencing projects, which make pathogen sequence data available as an epidemic unfolds [Gardy et al., 2015, Quick et al., 2016]. In general these new pathogen sequences arrive one at a time (or in small batches) into a background of existing sequences. Most phylogenetic inferences, however, are performed "from scratch" even when an inference has already been made on the previously available sequences. Thus projects such as `nextflu.org` [Neher and Bedford, 2015] incorporate new sequences into trees as they become available, but do so by recalculating the phylogeny from scratch at each update using a fast approximation to maximum likelihood inference, rather than a Bayesian method.

Modern researchers using phylogenetics are in the situation of having previous inferences, having new sequences, and yet having no principled method to incorporate those new sequences into existing inferences. Existing methods either treat a previous point estimate as an established fact and directly insert a new sequence into a phylogeny [Matsen et al., 2010, Berger et al., 2011], or use such a tree as a starting point for a new maximum-likelihood search [Izquierdo-Carrasco et al., 2014]. There is currently no method to update posterior distributions on phylogenetic trees with additional sequences.

In this paper we develop the theoretical foundations for an online Bayesian method for phylogenetic inference based on Sequential and Markov Chain Monte Carlo. Unlike previous applications of Sequential Monte Carlo (SMC) to phylogenetics [Bouchard-Côté et al., 2012, Bouchard-Côté, 2014, Wang et al., 2015], we develop and analyze algorithms that can update a posterior distribution as new sequence data becomes available. We first show a consistency result, demonstrating that the method samples from the correct distribution in the limit of a large number of particles in the SMC. Next we derive the first reported set of bounds on how phylogenetic likelihood surfaces change when new sequences are added. These bounds enable us to characterize the theoretical performance of sampling algorithms by developing a lower bound on the effective sample size (ESS) for a given number of particles. Surprisingly, this result holds even though the dimensions of the phylogenetic model grow with each new added sequence.

2 Mathematical setting

2.1 Background and notation

Throughout this paper, a *phylogenetic tree* (τ, l) is an unrooted tree τ with leaves labeled by a set of taxon names (e.g. species names), such that each edge e is associated with a non-negative number l_e. For each phylogenetic tree (τ, l), we will refer to τ as its *tree topology* and to l as the vector of *branch lengths*. We denote by $E(\tau)$ the set of all edges in trees with topology τ; any edge adjacent to a leaf is called a *pendant edge*, and any other edge is called an *internal edge*.

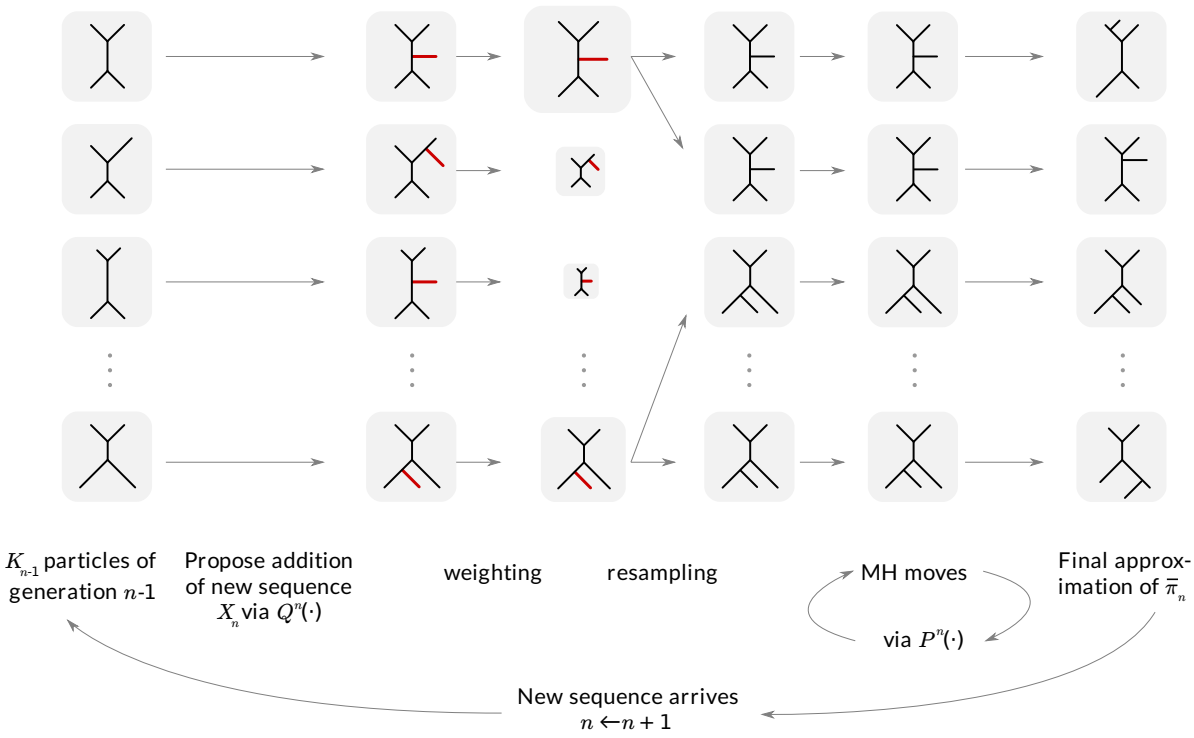

Figure 1: An overview of the Online Phylogenetic Sequential Monte Carlo algorithm.

We will employ the standard likelihood-based framework for statistical phylogenetics on discrete characters under the common assumption that alignment sites are IID [Felsenstein, 2004], which we now review briefly. Let Ω denote the set of character states and let $r = |\Omega|$. For DNA $\Omega = \{A, C, G, T\}$ and $r = 4$. We assume that the mutation events occur according to a continuous time Markov chain on states Ω with instantaneous rate matrix Ξ and stationary distribution ω. This rate matrix Ξ and the branch length l_e on the edge e define the transition matrix $G^e = e^{\Xi l_e}$ on edge e, where $G^e_{ij}(l_e)$ denotes the probability of mutating from state i to state j across the edge e (with length l_e).

In an online setting, the taxa $\{X_1, X_2, \ldots, X_N\}$ and their corresponding observed sequences $\{\psi_1, \psi_2, \ldots, \psi_N\}$, each of length S, arrive in a specific order, where N is a finite but large number. For all $n \leq N$, we consider the set of all phylogenetic trees that have $\{X_1, X_2, \ldots, X_n\}$ as their set of taxa and seek to sample from a sequence of probability distributions $\bar{\pi}_n$ of increasing dimension corresponding to phylogenetic likelihood functions [Felsenstein, 2004].

For a fixed phylogenetic tree (τ, l), the phylogenetic likelihood is defined as follows and will be denoted by $L(\tau, l)$. Given the set of observations $\psi(n) = (\psi_1, \psi_2, \ldots, \psi_S) \in \Omega^{n \times S}$ of length S up to time n, the likelihood of observing $\psi(n)$ given the tree has the form

$$L_n(\tau, l) = \prod_{u=1}^{S} \sum_{a^u} \omega(a^u_\rho) \prod_{(i,j) \in E(\tau)} G^{ij}_{a^u_i a^u_j}(l_{(i,j)})$$

where a^u ranges over all extensions of ψ to the internal nodes of the tree, a^u_i denotes the assigned state of node i by a^u, ρ denotes the root of the tree. Although we designate a root for notational convenience, the methods and results we discuss apply equally to unrooted trees.

Given a proper prior distribution with density π_0 imposed on branch lengths and tree topologies, the target posterior distributions can be computed as $\bar{\pi}_n(\tau, l) \sim L_n(\tau, l)\pi_0(\tau, l)$. We will also denote

by $\hat{\pi}_n(\tau, l)$ the un-normalized measure $L_n(\tau, l)\pi_0(\tau, l)$.

Throughout the paper, we assume that the phylogenetic trees of interest all have non-negative branch lengths bounded from above by $b > 0$ and denote by $\mathcal{T}_n$ the set of all such trees. To enable integration on tree spaces and define $\bar{\pi}_n$, we consider the natural probability measure on $\mathcal{T}_n$: the set $\mathcal{T}_n$ is viewed as the product space of the space of all possible tree topologies (with uniform measure) and the space of all branch lengths $[0, b]^{2n-3}$ (with Lebesgue measure). These measures can be written as

$$d\mu_n(\tau, l) = \frac{1}{V_n} d\nu_n(l) = \left(\frac{1}{V_n} d\tau\right) \left(\prod_{e \in E(\tau)} d\nu(l_e)\right),$$

where $V_n = (2n-3)!!$ is the number of different topologies of $\mathcal{T}_n$, l_e is the length of edge e, $d\tau$ is the counting measure on the set of all topologies on $\mathcal{T}_n$, and $d\nu$ is the Lebesgue measure on $\mathbb{R}^+$.

2.2 Sequential Monte Carlo

SMC methods are designed to approximate a sequence of probability distributions changing through time. These probability distributions may be of increasing dimension or complexity. They track the sequence of probability distributions of interest by producing a discrete representation of the distribution $\bar{\pi}_n$ at each iteration n through a random collection of weighted particles. After each generation, new sequences arrive and the collection of particles is updated to represent the next target distribution. While the details of the algorithms might vary, the main idea of SMC interspersed with MCMC sampling can be described as follows.

At the beginning of each iteration n, a list of K_n particles $p_1^n, \ldots, p_{K_n}^n$ are maintained along with a positive weight w_i^n associated with each particle p_i^n. These weighted particles form an un-normalized measure and a corresponding normalized empirical measure

$$\hat{\pi}_{n,K_n} = \sum_{i=1}^{K_{n+1}} w_i^n \delta_{p_i^n}(\cdot) \quad \text{and} \quad \bar{\pi}_{n,K_n} = K_n^{-1} \hat{\pi}_{n,K_n}$$

such that $\bar{\pi}_{n,K_n}$ approximates $\bar{\pi}_n$. A new list of K_{n+1} particles is then created in three steps: selection, Markov transition and mutation.

The aim of the selection step is to obtain an unweighted empirical distribution of the weighted measure $\hat{\pi}_{n,K_n}$ by discarding samples with small weights and allowing samples with large weights to reproduce. Formally, after selection we obtain the unweighted measure

$$\hat{\alpha}_{n,K_{n+1}} = \sum_{i=1}^{K_n} K_{n+1,i}\, \delta_{p_i^n}(\cdot)$$

where $K_{n+1,i}$ is the multiplicity of particle p_i^n, sampled from a multinomial distribution parameterized by the weights w_i^n. We denote the particles obtained after this step by s_i^n.

The scheme employed in the selection step introduces some Monte Carlo error. Moreover, when the distribution of the weights from the previous generation is skewed, the particles having high importance weights might be over-sampled. This results in a depletion of samples (or *path degeneracy*): after some generations, numerous particles are in fact sharing the same ancestor. A Markov transition step can be employed to alleviate this sampling bias, during which MCMC steps are run separately on each particle s_i^n for a certain amount of time to obtain a new independent sample m_i^n with (unweighted) measure denoted $\hat{\beta}_{n,K_{n+1}}$.

Finally, in the mutation step, new particles $t_1^{n+1}, \ldots, t_{K_{n+1}}^{n+1}$ are created from a proposal distribution Q^n and are weighted by an appropriate weight function h. If we assume further that for

each state t, there exists a unique state s, denoted by $\varrho(t)$, such that $Q^n(\varrho(t), t) > 0$, then h can be chosen as

$$h(t) = \frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))\ Q^n(\varrho(t), t)}. \tag{2.1}$$

The process is then iterated until $n = N$.

For convenience, we will denote the unnormalized empirical measures of the particles right after step n by $\hat{\alpha}_{n,K_{n+1}}$, $\hat{\beta}_{n,K_{n+1}}$ and $\hat{\lambda}_{n,K_{n+1}}$, respectively. Similarly, the corresponding normalized distributions will be denoted by $\bar{\alpha}_{n,K_{n+1}}$, $\bar{\beta}_{n,K_{n+1}}$ and $\bar{\lambda}_{n,K_{n+1}}$.

3 Online phylogenetic inference via Sequential Monte Carlo

Here we develop Online Phylogenetic sequential Monte Carlo (OPSMC) methods that continually update phylogenetic posteriors as new molecular sequences are added. In contrast to the traditional setting of SMC, for OPSMC when the number of leaves n of the particles increases, not only does the local dimension of the space $\mathcal{T}_n$ increase linearly, the number of different topologies in $\mathcal{T}_n$ also increases super-exponentially in n. Careful constructions of the proposal distribution Q^n, which will build $n+1$-taxon trees out of n-taxon trees, and the Markov transition kernel P^n are essential to cope with this increasing complexity.

Given two trees r and r' in the tree space $\mathcal{T} = \bigcup\ \mathcal{T}_n$, we say that r' *covers* r if there exists n such that $r \in \mathcal{T}_n$, $r' \in \mathcal{T}_{n+1}$, and r can be obtained from r' by removing the taxon X_{n+1} and its corresponding edge. This definition is analogous to the covering definition of Wang et al. [2015], although is distinct in the setting of online inference. The proposal distributions Q^n will be designed in such a way that the following criterion holds.

Criterion 3.1. *At every step of the OPSMC sampling process, the proposal density Q^n satisfies $Q^n(r, r') > 0$ if and only if r' covers r.*

Under this criterion, for every tree $t \in \mathcal{T}_{n+1}$, there exists a unique tree $\varrho(t)$ in $\mathcal{T}_n$ such that $Q^n(\varrho(t), t) > 0$ and thus a weight function of the form (2.1) can be used.

To obtain an $(n+1)$-taxon tree from an n-taxon tree, a proposal strategy Q^n must specify:

1. an edge e to which the new pendant edge is added,
2. the position x on that edge to attach the new pendant edge, and
3. the length y of the pendant edge.

The position x on an edge of a tree will be specified by its *distal* length, which is the distance from the attachment location to the end of the edge that is farthest away from the root of the tree. Different ways of choosing (e, x, y) lead to different sampling strategies and performances. Throughout the paper, we will investigate two different classes of sampling schemes: length-based proposals and likelihood-based proposals.

3.1 Length-based proposals

For length-based proposals:

1. the edge e is chosen from a multinomial distribution weighted by length of the edges,
2. the distal position x is selected from a distribution $P^e_X(x)$ across the edge length,

3. the pendant length y is sampled from a distribution $P_Y(y)$ with support contained in $[0, b]$.

For example, if these P distributions are uniform, we obtain a uniform (with respect to Lebesgue measure) prior on attachment locations across the tree. We assume that

Assumption 3.2. *The densities p_X^e of the distal position on edge e and p_Y of the pendant edge lengths are absolutely continuous with respect to the Lebesgue measure on $[0, l_e]$ and $[0, b]$, respectively. Moreover,*

$$\frac{1}{l_e^2}\int_0^{l_e} \frac{1}{p_X^e(x)}\, d\nu(x) \;\leq C \qquad \text{and} \qquad \int_0^{\infty} \frac{1}{p_Y(y)}\, d\nu(y) \;< \infty.$$

where l_e denotes the length of edge e and C is independent of l_e.

We note that, for any density function ψ on $[0, 1]$ such that $1/\psi$ is integrable, the family of proposals $\psi_l(x) = \frac{1}{l}\psi(\frac{x}{l})$ satisfies Assumption 3.2. The densities p_X^e and p_Y are assumed to be absolutely continuous to ensure Criterion 3.1 holds.

As we will discuss in later sections, to make sure that the proposals can capture the posterior distributions $\bar{\pi}_n$ efficiently, some regularity conditions on $\bar{\pi}_n$ are also necessary. These conditions are formalized in terms of a lower bound on the posterior expectation of $\zeta(s)$, the average branch length of s for a given tree $s \in \mathcal{T}_n$.

Assumption 3.3 (Assumption on the average branch length)**.** *There exist positive constants c (independent of n) such that for each n*

$$c \leq \int_{\mathcal{T}_n} \bar{\pi}_n(r)\zeta(r)\, dr$$

where $\zeta(r)$ denotes the average of branch lengths of the tree r.

3.2 Likelihood-based proposals

In the likelihood-based approach, the edge e (from the tree r) is chosen from a multinomial distribution weighted by a likelihood-based utility function $f(s, e)$. Similarly, the distributions $P_X^e(x)$ and $P_Y(y)$ might also be guided by information about the likelihood function. Likelihood-based proposals are capable of capturing the posterior distribution more efficiently, but with an additional cost for computing the likelihoods.

We define the average likelihood utility function

$$\mathcal{G}_n(r, e) = \int_{x,y} \hat{\pi}_{n+1}(T(r, e, x, y))\, dx\, dy$$

and use it as the prototype for likelihood-based utility functions. The likelihood-based utility function $f(r, e)$ is assumed to satisfy the following assumption.

Assumption 3.4. *There exist $c_1, c_2 > 0$ such that $c_1\mathcal{G}_n(r, e) \leq f_n(r, e) \leq c_2\mathcal{G}_n(r, e)$ for all r, e.*

The following lemma (proven in the Appendix) establishes that the maximum likelihood utility function also satisfies Assumption 3.4.

Lemma 3.5. *Let $f_n(r, e) = b\, l_e \sup_{x,y}\ \hat{\pi}_{n+1}(T(r, e, x, y))$, there exists $c_3 > 0$ independent of n such that $\mathcal{G}_n(r, e) \leq f_n(r, e) \leq c_3\mathcal{G}_n(r, e)$ for all s, e.*

As for the length-based proposal, we assume the following conditions on the distal position and pendant edge length proposals for the likelihood-based approach.

Assumption 3.6. *The densities p_X^e and p_Y are absolutely continuous with respect to the Lebesgue measure on $[0, l_e]$ and $[0, b]$, respectively. Moreover, there exists a_0 independent of n such that*

$$\sup_{x,y} \frac{1}{p_X^e(x)} \frac{1}{p_Y(y)} \leq a_0.$$

3.3 Markov transition kernels

Besides the SMC proposal strategy Q^n, it is also important to choose an appropriate Markov transition kernel P^n to have an effective OPSMC algorithm. It is worth noting that the problem of sample depletion is even more severe for OPSMC, since after each generation, the sampling space actually expands in dimensionality and complexity. To alleviate this sampling bias, MCMC steps are run separately on each particle for a certain amount of time to obtain new independent samples. We require the following criterion, which is as expected for any Markov transition kernel used in standard MCMC.

Criterion 3.7. *At every step of the OPSMC sampling process, the Markov transition kernel P^n has $\bar{\pi}_n$ as its invariant measure.*

As we will see later in the proof of consistency of OPSMC, Criterion 3.7 is the only assumption to be imposed on the Markov transition kernel. This leaves us with a great degree of freedom to improve the efficiency of the sampling algorithm without damaging its theoretical properties. For example, this allows us to use global information provided by the population of particles, such as effective sample size [Beskos et al., 2014], to guide the proposal, or to define a transition kernel on the whole set (or some subset) of particles [Andrieu et al., 2001]. In the context of phylogenetics, we can design a sampler that recognizes subtrees that have been insufficiently sampled, and samples more particles to improve the effective sample size within such regions. Similarly, one can use samplers that rearrange the tree structure in the neighborhood of newly added pendant edges.

4 Consistency of online phylogenetic SMC

In this section, we establish the consistency of OPSMC in the limit of a large number of particles by induction on the number of taxa n; that is, for every $n < N$, assuming that $\bar{\pi}_{n,K_n} \to \bar{\pi}_n$, we will prove that $\bar{\pi}_{n+1,K_{n+1}} \to \bar{\pi}_{n+1}$. We note that although the measures mentioned above are indexed by K_n, they implicitly depend on the number of particles from the previous generations. Thus, the convergence should be interpreted in the sense of when the number of particles of all generations approaches infinity.

The mode of convergence used in this section is "weak convergence", in which we say $\mu_K \to \mu$ if for every appropriate test function ϕ we have $\lim_{K\to\infty} \int \phi(t) d\mu_K(t) = \int \phi(t) d\mu(t)$. We will use $\mu(\phi)$ to denote $\int \phi(t) d\mu(t)$ for any measures μ and test functions ϕ.

For convenience, let L and K be the number of particles at n^{th} and $(n+1)^{\text{st}}$ generation, respectively. Recall that the normalized distributions after the substeps of OPSMC are denoted by $\bar{\alpha}_{n,K}$, $\bar{\beta}_{n,K}$ and $\bar{\lambda}_{n,K}$, we have the following lemma, proven in the Appendix.

Lemma 4.1. *Assume that Criteria 3.1 and 3.7 are satisfied. If we define*

$$\bar{\lambda}_n(t) := \bar{\pi}_n(\varrho(t))Q^n(\varrho(t), t) \qquad \text{and} \qquad h(t) := \frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))\ Q^n(\varrho(t), t)}$$

then the following statements hold.

1. *If $\bar{\pi}_{n,L}$ converges to $\bar{\pi}_n$, then $\bar{\alpha}_{n,K}$ converges to $\bar{\pi}_n$.*
2. *If $\bar{\alpha}_{n,K}$ converges to $\bar{\pi}_n$, then $\bar{\beta}_{n,K}$ converges to $\bar{\pi}_n$.*
3. *If $\bar{\beta}_{n,K}$ converges to $\bar{\pi}_n$, then $\bar{\lambda}_{n,K}$ converges to $\bar{\lambda}_n$.*
4. *$h(t)\bar{\lambda}_n(t)$ is proportional to $\bar{\pi}_{n+1}(t)$.*
5. *If $\bar{\lambda}_{n,K}$ converges to $\bar{\lambda}_n$, then $\bar{\pi}_{n+1,K}$ converges to $\bar{\pi}_{n+1}$.*

We note that when $n = 0$, the set of all rooted trees with no taxa consists of a single tree ρ. Thus, if we use this single tree as the ensemble of particles at $n = 0$, then $\bar{\pi}_{0,K_0}$ is precisely $\bar{\pi}_0$. Alternatively, we can start with $n = 1$ and use some ergodic MCMC methods to create an ensemble of particles with stationary distribution $\bar{\pi}_1$. In either case, an induction argument with Lemma 4.1 gives the main theorem:

Theorem 4.2 (Consistency)**.** *If Criteria 3.1 and 3.7 are satisfied and the sampler starts at $n = 0$ by a list consisting of a single rooted tree with no taxa, or at $n = 1$ with an ensemble of particles created by an ergodic MCMC method with stationary distribution $\bar{\pi}_1$, then*

$$\bar{\pi}_{n,K_n}(\phi) \to \bar{\pi}_n(\phi) \qquad \text{as} \qquad K_1, K_2, \ldots K_n \to \infty$$

for every integrable test function $\phi : \mathcal{T}_n \to \mathbb{R}$ and $n \leq N$.

5 Characterizing changes in the likelihood landscapes when new sequences arrive

Although the consistency of OPSMC is guaranteed and informative OPSMC samplers can be developed by changing the Markov transition kernels, its applicability is constrained by an implicit assumption: the distance between target distributions of consecutive generations are not too large. Since SMC methods are built upon the idea of recycling particles from one generation to explore the target distribution of the next generation, it is obvious that one would never be able to design an efficient SMC sampler if $\bar{\pi}_n$ and $\bar{\pi}_{n+1}$ are effectively orthogonal.

While a condition on minor changes in the target distributions may be easy to verify in some applications, it is not straightforward in the context of phylogenetic inference. A similar question on how the "optimal" trees (under some appropriate measure of optimality) change has been studied extensively in the field, with negative results for almost all regular measures of optimality [Heath et al., 2008, Cueto and Matsen, 2011]. To the best of our knowledge, no previous work has been done to investigate how phylogenetic likelihood landscapes change when new sequences arrive.

In this section, we will establish that under some minor regularity conditions on the distribution described in the previous sections, the relative changes between target distributions from consecutive generations are uniformly bounded. This result enables us to provide a lower bound on the effective sample size of OPSMC algorithms in the next section.

We denote by $T(r, e, x, y)$ the tree obtained by adding an edge of length y to edge e of the tree r at distal position x. Thus, any tree t can be represented by $t = (\varrho(t), e(t), x, y)$, where $e(t)$ is the edge on which the pendant edge containing the most recent taxon is attached.

Lemma 5.1 (Change of variables)**.** *The map* $(r, e, x, y) \to T(r, e, x, y)$ *is bijective. Moreover,*

$$d\mu_{n+1}(t) = \frac{V_n}{V_{n+1}} dx\ dy\ de\ d\mu_n(r)$$

where de *is the counting measure on the set of edges of an* n*-tree, and again* $V_n = (2n-3)!!$*.*

This result allows us to derive the following Lemma (detailed proof is provided in the Appendix).

Lemma 5.2. *Consider an arbitrary tree* $t \in \mathcal{T}_{n+1}$ *obtained from the parent tree* $\varrho(t)$ *by choosing edge* e*, distal position* x *and pendant length* y*. Denote*

$$M(y) = \max_{ij} G_{ij}(y), \qquad m(y) = \min_{ij} G_{ij}(y) \qquad \text{and} \qquad \mathcal{Z}_n = \int_{s \in \mathcal{T}_n} \bar{\pi}_n(s)\zeta(s)\ ds.$$

We have

$$\frac{\bar{\pi}_{n+1}(t)}{\bar{\pi}_n(\varrho(t))} \leq \frac{1}{\mathcal{Z}_n} \frac{M(y)^S}{\int_0^b m(y)^S\, dy}, \qquad \forall t \in \mathcal{T}_{n+1}.$$

Sketch of proof. By using the one-dimensional formulation of the phylogenetic likelihood function derived in [Dinh and Matsen, 2016], we can prove that

$$\frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))} = \frac{L_{n+1}(t)}{L_n(\varrho(t))} \leq M(y)^S, \qquad \forall t \in \mathcal{T}_{n+1}. \tag{5.1}$$

Similarly, we have $\hat{\pi}_{n+1}(t)/\hat{\pi}_n(\varrho(t)) \geq m(y)^S$ for all $t \in \mathcal{T}_{n+1}$.

Recall that $\zeta(r)$ is the average branch length of r. Using the fact that for a fixed tree r, $\int_0^{l_e} dx = l_e$ and $\sum_e l_e = (2n-3)\zeta(r)$, we have

$$\begin{aligned}
\|\hat{\pi}_{n+1}\| &= \int_{t \in \mathcal{T}_{n+1}} \hat{\pi}_{n+1}(t)\, dt \geq \int_{t \in \mathcal{T}_{n+1}} m(y)^S \hat{\pi}_n(\varrho(t))\, dt \\
&= \int_{r,e,x,y} m(y)^S \hat{\pi}_n(r)\, dx\, dy\, \frac{V_n}{V_{n+1}} de\, dr \\
&= \frac{(2n-3)V_n}{V_{n+1}} \left(\int_0^b m(y)^S\, dy \right) \int_{r \in \mathcal{T}_n} \hat{\pi}_n(r)\zeta(r)\, dr
\end{aligned}$$

which implies

$$\frac{\bar{\pi}_{n+1}(t)}{\bar{\pi}_n(\varrho(t))} = \frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))} \frac{\|\hat{\pi}_n\|}{\|\hat{\pi}_{n+1}\|} \leq \frac{1}{\mathcal{Z}_n} \frac{M(y)^S}{\int_0^b m(y)^S\ dy}, \qquad \forall t \in \mathcal{T}_{n+1}.$$

□

6 Effective sample sizes of online phylogenetic SMC

In this section, we are interested in the asymptotic behavior of OPSMC in the limit of large K_n, i.e. when the number of particles of the sampler approaches infinity. This asymptotic behavior is illustrated via estimates of the effective sample size of the sampler with large numbers of particles. We note that although there are several studies on the stability of SMC as the time step grows, most of them focus on cases where the sequence of target distributions have a common state space of fixed dimension [Del Moral, 1998, Douc and Moulines, 2008, Künsch, 2005, Oudjane and Rubenthaler,

2005, Del Moral et al., 2009, Beskos et al., 2014]. In general, establishing stability bounds for SMC requires imposing some conditions on the effect of data at any step k to the target distribution at step $n \geq k$ [Crisan and Doucet, 2002, Chopin, 2004, Doucet and Johansen, 2009]. Lemma 5.2 helps validate a condition of this type.

The effective sample size [Beskos et al., 2014] of the particles at step $n+1$ is computed as

$$\mathrm{ESS}_{n+1} = \frac{\left(\sum_{i=1}^{K} w_i^{n+1}\right)^2}{\sum_{i=1}^{K} (w_i^{n+1})^2}.$$

The following result, proven in the Appendix, enables us to estimate the asymptotic behavior of the sample's ESS in various settings.

Theorem 6.1. *In the limit as the number of particles approaches infinity, we have*

$$\lim_{K\to\infty} \frac{K_{n+1}}{ESS_{n+1}} = \int_{t\in\mathcal{T}_{n+1}} \frac{\bar{\pi}_{n+1}^2(t)}{\bar{\pi}_n(\varrho(t))\ Q^n(\varrho(t),t)}\ dt.$$

This asymptotic estimate and the results on likelihood landscapes from the previous section allow us to prove the following Theorem.

Theorem 6.2 (Effective sample size of OPSMC for likelihood-based proposals)**.** *If Assumptions 3.3, 3.6 and 3.4 hold, then there exists $\alpha > 0$ independent of n such that $ESS_n \geq \alpha K_n$. That is, the effective sample size of an OPSMC with likelihood-based proposals are bounded below by a constant multiple of the number of particles. Moreover, if Assumption 3.3 does not hold, the effective sample size of OPSMC algorithms decays at most linearly as the dimension increases.*

Proof of Theorem 6.2. Define $f_n(r) = \sum_e f_n(r,e)$, we have

$$\int_r f_n(r)\ ds \leq c_2 \int_r \sum_e \int_{x,y} \hat{\pi}_{n+1}(T(r,e,x,y))\ dx\ dy\ dr = c_2(2n-3)\|\hat{\pi}_{n+1}\|.$$

Since edge e is chosen from a multinomial distribution weighted by $f_n(r,e)$, given any tree $t \in \mathcal{T}_{n+1}$ obtained from the parent tree $\varrho(t)$, chosen edge $e(t)$, distal position x and pendant length y,

$$Q^n(\varrho(t),t) = \frac{V_{n+1}}{V_n} \frac{f(\varrho(t),e(t))}{f(\varrho(t))}\ p_X(x)\ p_Y(y).$$

By Lemma 5.2 and the fact that $M(y) \leq 1$, we have

$$\frac{\bar{\pi}_{n+1}(t)}{\bar{\pi}_n(\varrho(t))} \leq \frac{1}{\mathcal{Z}_n} \frac{M(y)^S}{\int_0^b m(y)^S\ dy} \leq \frac{1}{u_1 \mathcal{Z}_n},$$

where $u_1 = \int_0^b m(y)^S\ dy$ and $\mathcal{Z}_n$ are defined as in the proof of Lemma 6.3. Using Assumptions 3.3 and 3.6, (8.5), Lemma 5.1 and similar arguments as in the previous proof, we have

$$\begin{aligned}
&\int_{t\in\mathcal{T}_{n+1}} \frac{\bar{\pi}_{n+1}^2(t)}{\bar{\pi}_n(\varrho(t))\ Q^n(\varrho(t),t)}\ dt \\
&\leq a_0 \frac{1}{u_1\mathcal{Z}_n} \left(\frac{V_n}{V_{n+1}}\right)^2 \int_{r,e} \frac{f_n(r)}{f_n(r,e)} \int_{x,y} \bar{\pi}_{n+1}(T(r,e,x,y))\ dx\ dy\ dr\ de \\
&\leq \frac{a_0}{c_1} \frac{1}{u_1\mathcal{Z}_n} \frac{1}{\|\hat{\pi}_{n+1}\|} \left(\frac{V_n}{V_{n+1}}\right)^2 \left(\int_r f_n(r) ds\right) \left(\int_e de\right) \\
&\leq (2n-3)^2 a_0 \frac{c_2}{c_1} \frac{1}{u_1\mathcal{Z}_n} \left(\frac{V_n}{V_{n+1}}\right)^2 = a_0 \frac{c_2}{c_1} \frac{1}{u_1\mathcal{Z}_n}.
\end{aligned}$$

Thus by Theorem 6.1 there exists $\alpha > 0$ independent of K_n and n such that $\mathrm{ESS}_n \geq \alpha K_n$. We also note that without the assumption on average branch lengths, a crude estimate gives $\mathcal{Z}_n \geq \mathcal{Z}_1/n$, which leads to a linear decay in the upper bound on the ESS. □

We also have similar estimates for length-based proposals (see Appendix for proof):

Theorem 6.3 (Effective sample size of OPSMC for length-based proposals)**.** *If Assumptions 3.2 and 3.3 hold, then the effective sample size of OPSMC with length-based proposals are bounded below by a constant multiple of the number of particles. Moreover, if Assumption 3.3 does not hold, the effective sample size of OPSMC algorithms decays at most quadratically as the dimension increases.*

In summary, we are able to prove that in many settings, the effective sample size of OPSMC is bounded from below. These results are interesting, since in the general case it is known that SMC-type algorithms may suffer from the curse-of-dimensionality: when the dimension of the problem increases, the number of the particles must increase exponentially to maintain a constant effective sample size [Chopin, 2004, Bengtsson et al., 2008, Bickel et al., 2008, Snyder et al., 2008].

7 Discussion

In this paper, we establish foundations for Online Phylogenetic Sequential Monte Carlo (OPSMC), including essential theoretical convergence results. We prove that under some mild regularity conditions and with carefully constructed proposals, the OPSMC sampling algorithm is consistent. This includes relaxing the condition used in Bouchard-Côté et al. [2012], in which the authors assume that the weight of the particles are bounded from above. We then investigate two different classes of sampling schemes for online phylogenetic inference: length-based proposals and likelihood-based proposals. In both cases, we show the effective sample size to be bounded below by a multiple of the number of particles.

The consistency and convergence results in this paper apply to a variety of sampling strategies. One possibility would be for an algorithm to use a large number of particles, directly using the SMC machinery to approximate the posterior. Alternatively, the SMC part of the sampler could be quite limited, resulting in an algorithm which combines many independent parallel MCMC runs in a principled way. As described above, the SMC portion of the algorithm enables MCMC transition kernels that would normally be disallowed by the requirement of preserving detailed balance. For example, one could use a kernel that focuses effort around the part of the tree which has recently been disturbed by adding a new sequence.

In the future we will develop efficient and practical implementations of these ideas. Many challenges remain. For example, the exclusive focus of this paper has been on the tree structure, consisting of topology and branch lengths. However, Bayesian phylogenetics algorithms typically co-estimate mutation model parameters along with tree structures. Although proposals for other model parameters can be obtained by particle MCMC [Andrieu et al., 2010], we have not attempted to incorporate it into the current SMC framework. In addition, we note that the input for this type of phylogenetics algorithm consists of a *multiple sequence alignment* (MSA) of many sequences, rather than just individual sequences themselves. This raises the question of how to maintain an up-to-date MSA. Programs exist to add sequences into existing MSAs [Caporaso et al., 2010, Katoh and Standley, 2013], although from a statistical perspective, it could be preferable to jointly estimate a sequence alignment and tree posterior [Suchard and Redelings, 2006]. It is an open question how that could be done in an online fashion, although in principle it could be facilitated by some modifications to the sequence addition proposals described here.

References

Christophe Andrieu, Arnaud Doucet, and Elena Punskaya. Sequential Monte Carlo methods for optimal filtering. In *Sequential Monte Carlo Methods in Practice*, pages 79–95. Springer, 2001.

Christophe Andrieu, Arnaud Doucet, and Roman Holenstein. Particle markov chain monte carlo methods. *J. R. Stat. Soc. Series B Stat. Methodol.*, 72(3):269–342, 2010. ISSN 1369-7412. doi: 10.1111/j.1467-9868.2009.00736.x. URL http://dx.doi.org/10.1111/j.1467-9868.2009.00736.x.

Thomas Bengtsson, Peter Bickel, Bo Li, et al. Curse-of-dimensionality revisited: Collapse of the particle filter in very large scale systems. In *Probability and statistics: Essays in honor of David A. Freedman*, pages 316–334. Institute of Mathematical Statistics, 2008.

Simon A Berger, Denis Krompass, and Alexandros Stamatakis. Performance, accuracy, and web server for evolutionary placement of short sequence reads under maximum likelihood. *Syst. Biol.*, 60(3):291–302, May 2011. ISSN 1063-5157, 1076-836X. doi: 10.1093/sysbio/syr010. URL http://dx.doi.org/10.1093/sysbio/syr010.

Alexandros Beskos, Dan Crisan, Ajay Jasra, et al. On the stability of sequential monte carlo methods in high dimensions. *The Annals of Applied Probability*, 24(4):1396–1445, 2014.

Peter Bickel, Bo Li, Thomas Bengtsson, et al. Sharp failure rates for the bootstrap particle filter in high dimensions. In *Pushing the limits of contemporary statistics: Contributions in honor of Jayanta K. Ghosh*, pages 318–329. Institute of Mathematical Statistics, 2008.

Alexandre Bouchard-Côté. SMC (sequential monte carlo) for bayesian phylogenetics. In Ming-Hui Chen, Lynn Kuo, and Paul O Lewis, editors, *Bayesian Phylogenetics: Methods, Algorithms, and Applications*. CRC Press, 2014.

Alexandre Bouchard-Côté, Sriram Sankararaman, and Michael I Jordan. Phylogenetic inference via sequential monte carlo. *Systematic biology*, 61(4):579–593, 2012.

J Gregory Caporaso, Kyle Bittinger, Frederic D Bushman, Todd Z DeSantis, Gary L Andersen, and Rob Knight. PyNAST: a flexible tool for aligning sequences to a template alignment. *Bioinformatics*, 26(2):266–267, 15 January 2010. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/btp636. URL http://dx.doi.org/10.1093/bioinformatics/btp636.

Nicolas Chopin. Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *Annals of Statistics*, pages 2385–2411, 2004.

Dan Crisan and Arnaud Doucet. A survey of convergence results on particle filtering methods for practitioners. *IEEE Transactions on signal processing*, 50(3):736–746, 2002.

María Angélica Cueto and Frederick A Matsen. Polyhedral geometry of phylogenetic rogue taxa. *Bulletin of Mathematical Biology*, 73(6):1202–1226, 2011.

Pierre Del Moral. A uniform convergence theorem for the numerical solving of the nonlinear filtering problem. *Journal of Applied Probability*, pages 873–884, 1998.

Pierre Del Moral, Frédéric Patras, Sylvain Rubenthaler, et al. Tree based functional expansions for Feynman–Kac particle models. *The Annals of Applied Probability*, 19(2):778–825, 2009.

Vu Dinh and Frederick A Matsen. The shape of the one-dimensional phylogenetic likelihood function. *in press, The Annals of Applied Probability*, 2016. http://arxiv.org/abs/1507.03647.

Randal Douc and Eric Moulines. Limit theorems for weighted samples with applications to sequential Monte Carlo methods. *The Annals of Statistics*, pages 2344–2376, 2008.

Arnaud Doucet and Adam M Johansen. A tutorial on particle filtering and smoothing: Fifteen years later. 2009.

Joseph Felsenstein. *Inferring phylogenies*, volume 2. Sinauer Associates Sunderland, 2004.

Jennifer Gardy, Nicholas J Loman, and Andrew Rambaut. Real-time digital pathogen surveillance — the time is now. *Genome Biol.*, 16(1):155, 30 July 2015. ISSN 1465-6906. doi: 10.1186/s13059-015-0726-x. URL http://www.genomebiology.com/content/pdf/s13059-015-0726-x.pdf.

Tracy A Heath, Shannon M Hedtke, and David M Hillis. Taxon sampling and the accuracy of phylogenetic analyses. *Journal of Systematics and Evolution*, 46(3):239–257, 2008.

Fernando Izquierdo-Carrasco, John Cazes, Stephen A Smith, and Alexandros Stamatakis. PUmPER: phylogenies updated perpetually. *Bioinformatics*, 30(10):1476–1477, 15 May 2014. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/btu053. URL http://dx.doi.org/10.1093/bioinformatics/btu053.

Kazutaka Katoh and Daron M Standley. MAFFT multiple sequence alignment software version 7: improvements in performance and usability. *Mol. Biol. Evol.*, 30(4):772–780, April 2013. ISSN 0737-4038, 1537-1719. doi: 10.1093/molbev/mst010. URL http://dx.doi.org/10.1093/molbev/mst010.

Hans R Künsch. Recursive Monte Carlo filters: algorithms and theoretical analysis. *Annals of Statistics*, pages 1983–2021, 2005.

Frederick Matsen, Robin Kodner, and E Virginia Armbrust. pplacer: linear time maximum-likelihood and bayesian phylogenetic placement of sequences onto a fixed reference tree. *BMC Bioinformatics*, 11(1):538, 2010. ISSN 1471-2105. doi: 10.1186/1471-2105-11-538. URL http://www.biomedcentral.com/1471-2105/11/538.

Richard A Neher and Trevor Bedford. nextflu: Real-time tracking of seasonal influenza virus evolution in humans. *Bioinformatics*, 26 June 2015. ISSN 1367-4803, 1367-4811. doi: 10.1093/bioinformatics/btv381. URL http://dx.doi.org/10.1093/bioinformatics/btv381.

Nadia Oudjane and Sylvain Rubenthaler. Stability and uniform particle approximation of nonlinear filters in case of non ergodic signals. *Stochastic Analysis and Applications*, 23(3):421–448, 2005.

Joshua Quick, Nicholas J Loman, Sophie Duraffour, Jared T Simpson, Ettore Severi, Lauren Cowley, Joseph Akoi Bore, Raymond Koundouno, Gytis Dudas, Amy Mikhail, Nobila Ouédraogo, Babak Afrough, Amadou Bah, Jonathan H J Baum, Beate Becker-Ziaja, Jan Peter Boettcher, Mar Cabeza-Cabrerizo, Álvaro Camino-Sánchez, Lisa L Carter, Juliane Doerrbecker, Theresa Enkirch, Isabel García-Dorival, Nicole Hetzelt, Julia Hinzmann, Tobias Holm, Liana Eleni Kafetzopoulou, Michel Koropogui, Abigael Kosgey, Eeva Kuisma, Christopher H Logue, Antonio Mazzarelli, Sarah Meisel, Marc Mertens, Janine Michel, Didier Ngabo, Katja Nitzsche, Elisa Pallasch, Livia Victoria Patrono, Jasmine Portmann, Johanna Gabriella Repits, Natasha Y

Rickett, Andreas Sachse, Katrin Singethan, Inês Vitoriano, Rahel L Yemanaberhan, Elsa G Zekeng, Trina Racine, Alexander Bello, Amadou Alpha Sall, Ousmane Faye, Oumar Faye, N'faly Magassouba, Cecelia V Williams, Victoria Amburgey, Linda Winona, Emily Davis, Jon Gerlach, Frank Washington, Vanessa Monteil, Marine Jourdain, Marion Bererd, Alimou Camara, Hermann Somlare, Abdoulaye Camara, Marianne Gerard, Guillaume Bado, Bernard Baillet, Déborah Delaune, Koumpingnin Yacouba Nebie, Abdoulaye Diarra, Yacouba Savane, Raymond Bernard Pallawo, Giovanna Jaramillo Gutierrez, Natacha Milhano, Isabelle Roger, Christopher J Williams, Facinet Yattara, Kuiama Lewandowski, James Taylor, Phillip Rachwal, Daniel J Turner, Georgios Pollakis, Julian A Hiscox, David A Matthews, Matthew K O'Shea, Andrew Mcd Johnston, Duncan Wilson, Emma Hutley, Erasmus Smit, Antonino Di Caro, Roman Wölfel, Kilian Stoecker, Erna Fleischmann, Martin Gabriel, Simon A Weller, Lamine Koivogui, Boubacar Diallo, Sakoba Keïta, Andrew Rambaut, Pierre Formenty, Stephan Günther, and Miles W Carroll. Real-time, portable genome sequencing for ebola surveillance. *Nature*, 530 (7589):228–232, 11 February 2016. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature16996. URL http://dx.doi.org/10.1038/nature16996.

Chris Snyder, Thomas Bengtsson, Peter Bickel, and Jeff Anderson. Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136(12):4629–4640, 2008.

Marc A Suchard and Benjamin D Redelings. BAli-Phy: simultaneous bayesian inference of alignment and phylogeny. *Bioinformatics*, 22(16):2047–2048, 15 August 2006. ISSN 1367-4803. doi: 10.1093/bioinformatics/btl175. URL http://bioinformatics.oxfordjournals.org/content/22/16/2047.abstract.

Liangliang Wang, Alexandre Bouchard-Côté, and Arnaud Doucet. Bayesian phylogenetic inference using a combinatorial sequential monte carlo method. *J. Am. Stat. Assoc.*, 110(512):1362–1374, 2015. ISSN 0162-1459. doi: 10.1080/01621459.2015.1054487. URL http://dx.doi.org/10.1080/01621459.2015.1054487.

8 Appendix

Proof of Lemma 3.5. The lower bound is straightforward. For the upper bound, consider $(x, y) \in [0, l_e] \times [0, b]$ and fix $\delta > 0$; by the same arguments as in the proof of Lemma 5.2, we have

$$\hat{\pi}_{n+1}(T(r, e, x, y)) \geq m(\delta)^S \hat{\pi}_n(r) \qquad \forall y \geq \delta.$$

Thus, if we define

$$A = \{(x, y) \in [0, l_e] \times [0, b] : \hat{\pi}_{n+1}(T(r, e, x, y)) \geq m(\delta)^S \hat{\pi}_n(r)\},$$

then we have $|A| \geq (b - \delta) l_e$ and

$$\begin{aligned} \mathcal{G}_n(r, e) &= \int_{x,y} \hat{\pi}_{n+1}(T(r, e, x, y))\ dx\ dy \\ &\geq \int_A \hat{\pi}_{n+1}(T(r, e, x, y))\ dx\ dy \geq (b - \delta) l_e\ m(\delta)^S \hat{\pi}_n(r). \end{aligned}$$

On the other hand, from Lemma 5.2, we have $f_n(r, e) \leq b\, l_e\ \hat{\pi}_n(r) M(b)^S$.

By choosing $\delta = b/2$, we obtain

$$f_n(r, e) \leq 2 \frac{M(b)^S}{m(b/2)^S}\ \mathcal{G}_n(r, e)$$

which completes the proof. □

Proof of Lemma 4.1. (1). Assume that $\bar{\pi}_{n,L}$ converges to $\bar{\pi}_n$.

$$|\bar{\alpha}_{n,K}(\phi) - \bar{\pi}_n(\phi)| \leq \left| \frac{1}{K} \sum_{i=1}^{L} K_{n+1,i}\ \phi(p_i^n) - \frac{1}{\|w\|} \sum_{i=1}^{L} w_i\ \phi(p_i^n) \right| + |\bar{\pi}_{n,L}(\phi) - \bar{\pi}_n(\phi)|.$$

By the strong law of large numbers,

$$\limsup_{K \to \infty} |\bar{\alpha}_{n,K}(\phi) - \bar{\pi}_n(\phi)| \leq |\bar{\pi}_{n,L}(\phi) - \bar{\pi}_n(\phi)|$$

This implies that when $K, L \to \infty$, we have $\bar{\alpha}_{n,K}(\phi) \to \bar{\pi}_n(\phi)$.

(2). The rationale behind the use of MCMC moves is based on the observation that if the unweighted particles are distributed according to $\bar{\pi}_n$, then when we apply a Markov transition kernel P of invariant distribution $\bar{\pi}_n$ to any particle, the new particles are still distributed according to the posterior distribution of interest.

Formally, if $\bar{\alpha}_{n,K}(\phi) \to \pi_n(\phi)$ for every integrable test function $\phi : \mathcal{T}_n \to \mathbb{R}$, by choosing $\phi = P^r(\cdot, A)$ for any measurable set $A \subset \mathcal{T}_{n+1}$, we deduce that

$$\bar{\beta}_{n+1,K}(A) = \sum_{i=1}^{K} P^r(s, A)\ \bar{\alpha}_{n,K}(s_i^n) \xrightarrow{K \to \infty} \int_{\mathcal{T}_n} P^r(s, A)\ \bar{\pi}_n(s)\ ds = \bar{\pi}_n(A)$$

since the Markov kernel P is invariant with respect to $\hat{\pi}_n$. Therefore, for any measurable function $\phi : \mathcal{T}_n \to \mathbb{R}$, we have that $\bar{\beta}_{n,K}(\phi)$ converges to $\bar{\pi}_n(\phi)$.

(3). Since $\bar{\beta}_{n,K}(\phi) \to \pi_n(\phi)$ for every integrable test function $\phi : \mathcal{T}_n \to \mathbb{R}$, by choosing $\phi = Q^n(\cdot, t)$ for all $t \in \mathcal{T}_{n+1}$, we deduce that

$$\begin{aligned}\bar{\lambda}_{n,K}(t) &= \sum_{i=1}^{K} Q^n(m_i^n, t)\bar{\beta}_{n,K}(m_i^n).\\ &\xrightarrow{K\to\infty} \int_{\mathcal{T}_n} Q^n(m,t)\bar{\pi}_n(m)\, dm = \bar{\pi}_n(\varrho(t))\, Q^n(\varrho(t),t) = \bar{\lambda}_n(t).\end{aligned}$$

Moreover, for every measurable set $A \subset \mathcal{T}_{n+1}$, we can use the same argument to prove that

$$\begin{aligned}\bar{\lambda}_{n,K}(A) &= \sum_{i=1}^{K} Q^n(m_i^n, A)\bar{\beta}_{n,K}(m_i^n)\\ &\xrightarrow{K\to\infty} \int_{\mathcal{T}_n} Q^n(m,A)\bar{\pi}_n(dm) = \int_A \bar{\pi}_n(\varrho(t))Q^n(\varrho(t),t)\, dt = \bar{\lambda}_n(A).\end{aligned}$$

Therefore, for any measurable function $\phi : \mathcal{T}_{n+1} \to \mathbb{R}$, we also have $\bar{\lambda}_{n,K}(\phi)$ converges to $\bar{\lambda}_n(\phi)$.

(4) We note that

$$\begin{aligned}h(t)\bar{\lambda}_n(t) &= h(t)\,\bar{\pi}_n(\varrho(t))\, Q^n(\varrho(t),t)\\ &= \frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))\, Q^n(\varrho(t),t)}\,\frac{1}{\|\hat{\pi}_n\|}\,\hat{\pi}_n(\varrho(t))\, Q^n(\varrho(t),t)\\ &= \frac{1}{\|\hat{\pi}_n\|}\hat{\pi}_{n+1}(t).\end{aligned} \tag{8.1}$$

(5). Since the proposal Q^n and the Markov kernel P are assumed to be normalized, we have $\|\hat{\lambda}_{n,K}\| = \|\hat{\beta}_{n,K}\| = \|\hat{\alpha}_{n,K}\| = K$.

We have:

$$\begin{aligned}\frac{1}{K}\|\hat{\pi}_{n+1,K}\| &= \frac{1}{\|\hat{\lambda}_{n,K}\|}\sum_{i=1}^{K}\hat{\pi}_{n+1,K}(t_i^{n+1}) = \sum_{i=1}^{K} h(t_i^{n+1})\bar{\lambda}_{n,K}(t_i^{n+1})\\ &\xrightarrow{K\to\infty} \int_{\mathcal{T}_{n+1}} h(t)\bar{\lambda}_n(t)\, dt = \frac{1}{\|\hat{\pi}_n\|}\int_{\mathcal{T}_{n+1}}\hat{\pi}_{n+1}(t)\, dt = \frac{\|\hat{\pi}_{n+1}\|}{\|\hat{\pi}_n\|}.\end{aligned} \tag{8.2}$$

By a similar argument, we have

$$\begin{aligned}\bar{\pi}_{n+1,K}(\phi) = \frac{\frac{1}{K}\sum_{i=1}^{K}\phi(t_i^{n+1})\hat{\pi}_{n+1,K}(t_i^{n+1})}{\frac{1}{K}\|\hat{\pi}_{n+1,K}\|} &= \frac{\sum_{i=1}^{K}\phi(t_i^{n+1})h(t_i^{n+1})\bar{\lambda}_{n,K}(t_i^{n+1})}{\frac{1}{K}\|\hat{\pi}_{n+1,K}\|}\\ \xrightarrow{K\to\infty} \frac{\|\hat{\pi}_n\|}{\|\hat{\pi}_{n+1}\|}\int_{\mathcal{T}_n}\phi(t)h(t)\bar{\lambda}_n(t)dt &= \frac{\|\hat{\pi}_n\|}{\|\hat{\pi}_{n+1}\|}\int_{\mathcal{T}_n}\phi(t)\frac{\hat{\pi}_{n+1}(t)}{\|\hat{\pi}_n\|}dt = \bar{\pi}_{n+1}(\phi).\end{aligned}$$

In other words, $\bar{\pi}_{n+1,K}$ converges to $\bar{\pi}_{n+1}$. □

Proof of Theorem 6.1. By definition, we have

$$h(t) = \frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))\; Q^n(\varrho(t),t)}, \qquad \text{and} \qquad w_i^{n+1} = \hat{\pi}_{n+1,K_{n+1}}(t_i^{n+1}) = h(t_i^{n+1}).$$

Thus,

$$\begin{aligned}\bar{\pi}_{n+1,K_{n+1}}(h) &= \frac{\sum_{i=1}^{K_{n+1}} h(t_i^{n+1})\hat{\pi}_{n+1,K}(t_i^{n+1})}{\|\hat{\pi}_{n+1,K_{n+1}}\|} \\ &= \frac{\sum_{i=1}^{K_{n+1}} (w_i^{n+1})^2}{\sum_{i=1}^{K_{n+1}} w_i^{n+1}} = \frac{\sum_{i=1}^{K_{n+1}} w_i^{n+1}}{\mathrm{ESS}_{n+1}} = \frac{\|\hat{\pi}_{n+1,K_{n+1}}\|}{\mathrm{ESS}_{n+1}}.\end{aligned}$$

On the other hand, by applying Theorem 4.2 for $\phi \equiv h$, we have

$$\bar{\pi}_{n+1,K_{n+1}}(h) \quad \to \quad \bar{\pi}_{n+1}(h) = \frac{\|\hat{\pi}_{n+1}\|}{\|\hat{\pi}_n\|} \int_{t\in\mathcal{T}_{n+1}} \frac{\bar{\pi}_{n+1}^2(t)}{\bar{\pi}_n(\varrho(t))\ Q^n(\varrho(t),t)}\, dt$$

which completes the proof via the convergence result (8.2). □

Proof of Lemma 5.2. Let l_e be the length of the edge e and $G(\alpha)$ be the transition matrix across an edge of length α and k_u the observed value at site u of the newly added taxon. We follow the formulation of one-dimensional phylogenetic likelihood function as in [Dinh and Matsen, 2016] to fix all parameters except l_e and consider the likelihood of $\varrho(t)$ a function of l_e, we have

$$L_n(\varrho(t)) = \prod_{u=1}^{S} \left(\sum_{ij} d_{ij}^u G_{ij}^e(l_e) \right)$$

where d_{ij}^u the probability of observing i and j at the left and right nodes of e at the site index u, respectively (note that in [Dinh and Matsen, 2016] it is called d_{ij}^u). Similarly, by representing the likelihood of the tree t in terms of x, y and l_e, we have

$$L_{n+1}(t) = \prod_{u=1}^{S} \left(\sum_{ij} d_{ij}^u G_{ij}^e(l_e,x,y) \right) = \prod_{u=1}^{S} \left(\sum_{ij} d_{ij}^u \sum_m G_{im}(x) G_{mj}(l_e - x) G_{mk_u}(y) \right) \tag{8.3}$$

where the indices i, j, m are looped over all possible state characters. Since $G_{mk_u}(y) \le M(y)$ for all m and k_u, we deduce that

$$\frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))} = \frac{L_{n+1}(t)}{L_n(\varrho(t))} \le M(y)^S, \qquad \forall t \in \mathcal{T}_{n+1}. \tag{8.4}$$

Similarly, we have $\hat{\pi}_{n+1}(t)/\hat{\pi}_n(\varrho(t)) \ge m(y)^S$ for all $t \in \mathcal{T}_{n+1}$.

Recall that $\zeta(r)$ is the average branch length of r. Using the fact that for a fixed tree r, $\int_0^{l_e} dx = l_e$ and $\sum_e l_e = (2n-3)\zeta(r)$, we have

$$\begin{aligned}\|\hat{\pi}_{n+1}\| &= \int_{t\in\mathcal{T}_{n+1}} \hat{\pi}_{n+1}(t)\, dt \ge \int_{t\in\mathcal{T}_{n+1}} m(y)^S \hat{\pi}_n(\varrho(t))\, dt \\ &= \int_{r,e,x,y} m(y)^S \hat{\pi}_n(r)\, dx\, dy\, \frac{V_n}{V_{n+1}} de\, dr \\ &= \frac{(2n-3)V_n}{V_{n+1}} \left(\int_0^b m(y)^S\, dy \right) \int_{r\in\mathcal{T}_n} \hat{\pi}_n(r)\zeta(r)\, dr.\end{aligned}$$

Noting that $V_{n+1} = (2n-3)V_n$, we obtain

$$\frac{\|\hat{\pi}_{n+1}\|}{\|\hat{\pi}_n\|} \ge \left(\int_0^b m(y)^S\, dy \right) \mathcal{Z}_n \tag{8.5}$$

which implies

$$\frac{\bar{\pi}_{n+1}(t)}{\bar{\pi}_n(\varrho(t))} = \frac{\hat{\pi}_{n+1}(t)}{\hat{\pi}_n(\varrho(t))}\frac{\|\hat{\pi}_n\|}{\|\hat{\pi}_{n+1}\|} \leq \frac{1}{\mathcal{Z}_n}\frac{M(y)^S}{\int_0^b m(y)^S\ dy}, \qquad \forall t \in \mathcal{T}_{n+1}.$$

□

Proof of Theorem 6.3. Since the edge e is chosen from a multinomial distribution weighted by length of the edges, then given any tree $t \in \mathcal{T}_{n+1}$ obtained from the parent tree $\varrho(t)$ by choosing edge e, distal position x and pendant length y, we have

$$Q^n(\varrho(t), t) = \frac{V_{n+1}}{V_n}\frac{l_e(\varrho(t))}{l(\varrho(t))}\ p_X(x)\ p_Y(y)$$

where $l_e(r), l(r)$ are the length of edge e and the total tree length, respectively and V_n and V_{n+1} are the numbers of tree topologies of $\mathcal{T}_n$ and $\mathcal{T}_{n+1}$.

We denote

$$u_1 = \int_0^b m(y)^S\ dy, \qquad u_2 = \int_0^b \frac{M(y)^{2S}}{p_Y(y)}\ dy,$$

and recall that

$$\sum_e \frac{1}{l_e(r)}\int_0^{l_e(r)} \frac{1}{p_X^e(x)}\ dx\ \leq\ C\sum_e l_e(r)\ =\ Cl(r),$$

where C is the constant from Assumption 3.2, and

$$\mathcal{Z}_n = \int_{s\in\mathcal{T}_n} \bar{\pi}_n(r)\zeta(r)\ dr \geq c$$

from the assumption on the average branch length (Assumption 3.3). We have

$$\begin{aligned}
&\int_{t\in\mathcal{T}_{n+1}} \frac{\bar{\pi}_{n+1}^2(t)}{\bar{\pi}_n(\varrho(t))\ Q^n(\varrho(t),t)}\ dt = \int_{t\in\mathcal{T}_{n+1}} \frac{\bar{\pi}_{n+1}^2(t)}{\bar{\pi}_n^2(\varrho(t))}\frac{1}{Q^n(\varrho(t),t)}\ \bar{\pi}_n(\varrho(t))\ dt \\
&\leq \left(\frac{V_n}{V_{n+1}}\right)^2 \frac{1}{\mathcal{Z}_n^2}\int_{r,e,x,y} \frac{M(y)^{2S}}{u_1^2}\frac{l(r)}{l_e(r)}\frac{1}{p_X^e(x)}\frac{1}{p_Y(y)}\bar{\pi}_n(r)\ dx\ dy\ de\ dr \\
&= \left(\frac{V_n}{V_{n+1}}\right)^2 \frac{1}{\mathcal{Z}_n^2}\frac{u_2}{u_1^2}\int_{\mathcal{T}_n}\left(\sum_e \frac{1}{l_e(r)}\int_0^{l_e(r)}\frac{1}{p_X^e(x)}\ dx\right)\ l(r)\ \bar{\pi}_n(r)\ dr.
\end{aligned}$$

By the assumption of maximum branch length b, we have

$$\begin{aligned}
\int_{t\in\mathcal{T}_{n+1}} \frac{\bar{\pi}_{n+1}^2(t)}{\bar{\pi}_n(\varrho(t))\ Q^n(\varrho(t),t)}\ dt &\leq C(2n-3)^2\left(\frac{V_n}{V_{n+1}}\right)^2\frac{1}{\mathcal{Z}_n^2}\int_{\mathcal{T}_n}\bar{\pi}_n(r)\zeta^2(r)\ dr \\
&\leq \frac{Cb^2}{c^2}.
\end{aligned}$$

Thus by Theorem 6.1 there exists $\alpha > 0$ independent of K_n and n such that $\mathrm{ESS}_n \geq \alpha K_n$. □